Bayesian estimators of the Gamma distribution

A. Llera, C. F. Beckmann.

Technical report
Radboud University Nijmegen
Donders Institute for Brain Cognition and Behaviour.

arXiv:1607.03302v1 [stat.ME] 12 Jul 2016

Abstract

In this paper we introduce two Bayesian estimators for learning the parameters of the Gamma distribution. The first algorithm uses a well known unnormalized conjugate prior for the Gamma shape and the second one uses a non-linear approximation to the likelihood and a prior on the shape that is conjugate to the approximated likelihood. In both cases use the Laplace approximation to compute the required expectations. We perform a theoretical comparison between maximum likelihood and the presented Bayesian algorithms that allow us to provide non-informative parameter values for the priors hyper parameters. We also provide a numerical comparison using synthetic data. The introduction of these novel Bayesian estimators open the possibility of including Gamma distributions into more complex Bayesian structures, e.g. variational Bayesian mixture models.

1 Introduction

The Gamma distribution is the most considered when modeling positive data [1, 11, 9, 6]. The fastest and oldest method used to estimate the parameters of a Gamma distribution is the Method of Moments (MM) [1]. This method can be applied to the Gamma distribution since there exists an unique relationship between its first two moments and its parameters. The MM uses a Gaussian approximation to the moments to provide a closed form parameters estimation. Obviously this method is exact for the Gaussian distribution and less accurate for distributions deviating from Gaussianity. Another classic approach for parameter estimation is the well known maximum likelihood (ML), based in the maximization of the data log-likelihood; for the Gamma distribution there exist a closed form ML derivation for its scale parameter but not for the shape parameter. Two interesting algorithms to deal with this issue can be found in [9]. One of them uses a linear approximation to the otherwise intractable terms appearing in the derivative of the log-likelihood and the second one uses a more advanced non-linear approximation to the log-likelihood to obtain higher convergence rate [8, 9]. While the MM and ML approaches provide point estimates for the distribution parameter values, a full Bayesian estimation approach introduces distributions over the parameters [7]; when choosing conjugate prior distributions over the parameters [7, 2], the posterior expectations can be analytically derived. The conjugate prior for the Gamma rate parameter is known to be Gamma distributed but there exist no proper conjugate prior for the shape parameter.

In section 2 we present the Bayesian methods proposed in this article. In subsection 2.1 we review the conjugate prior for the rate (inverse scale) parameter. Since for the shape parameter there exist no proper conjugate prior distribution, here we propose two different strategies to deal

with this issue. In subsection 2.2 we consider the use of an unnormalized conjugate prior [4] and in section 2.3 we introduce an unnormalized prior on the shape which is conjugate with the accurate non-linear likelihood approximation presented in [8]. In both cases we use the Laplace approximation to compute the required expectations. In section 3.1 we provide a theoretical comparison between the two ML algorithms presented in [9] and the two presented Bayesian strategies. In section 3.2 we present a numerical comparison between the four algorithms. In section 3.3 we briefly consider the hyper parameters value choices for the Bayesian algorithms. We conclude the paper with a brief discussion in section 4. For completeness of the methodology and ease of reading, in Appendix A we review the Method of Moments estimation, and in Appendix B the two Maximum Likelihood algorithms introduced in [9].

2 Bayesian Learning (BL) for the Gamma distribution

Consider the Gamma distribution defined over the support $x \in \mathbb{R}^+$

$$\mathcal{G}(x|\alpha,\beta) = \frac{x^{\alpha-1}}{\Gamma(\alpha)\beta^{\alpha}} \exp(\frac{-x}{\beta}), \tag{1}$$

for any positive real parameter values α and β, denoting the distribution shape and scale respectively, and where Γ is the Gamma function. Given a vector of positive real values $\mathbf{x} = \{x_1, \ldots, x_n\}$, in the following subsections we present two different Bayesian strategies to find posterior estimations $\hat{\alpha}$ and $\hat{\beta}$ for the distribution parameters α and β.

To find the posterior probability of the Gamma parameters, $\theta = \{\alpha, \beta\}$, we use Bayes rule

$$p(\theta|\mathbf{x}) = \frac{p(\mathbf{x}|\theta)p(\theta)}{p(\mathbf{x})}, \tag{2}$$

where $\mathbf{x} = \{x_1, \ldots, x_n\}$ is some positive vector of observations. Since the denominator only depends on data the posterior is proportional to the likelihood multiplied by the prior

$$p(\theta|\mathbf{x}) \propto p(\mathbf{x}|\theta)p(\theta). \tag{3}$$

Obtaining analytical solutions for the parameters θ requires the use of conjugate priors [2, 4]. A prior is called conjugate with a likelihood function if the prior functional form remains unchanged after multiplication by the likelihood.

2.1 Rate parameter $R = \frac{1}{\beta}$

A well known conjugate prior for $R = \frac{1}{\beta}$, the rate parameter of the Gamma distribution, is a Gamma distribution parametrized using shape d and rate e [4],

$$p(R) = \mathcal{G}(R|d,e). \tag{4}$$

Given the observations vector $\mathbf{x}$, and multiplying its Gamma likelihood by the prior on the rate (4) we get its posterior, $q(R) = \mathcal{G}(R|\hat{d},\hat{e})$ with

$$\hat{d} = d + n\alpha, \qquad \hat{e} = e + \sum_{i=1}^{n} x_i. \tag{5}$$

Since R is Gamma distributed its posterior expectation is

$$\hat{R} = \frac{\hat{d}}{\hat{e}}. \tag{6}$$

Since the scale is the inverse of rate we have that $\hat{\beta} = \frac{\hat{e}}{\hat{d}}$.

2.2 Unnormalized prior on shape parameter α (BL1)

The unnormalized prior used for the shape parameter of the gamma distribution [4] has the form

$$p(\alpha) \propto \frac{a^{\alpha-1} R^{\alpha c}}{\Gamma(\alpha)^b}, \tag{7}$$

where R is the Gamma rate parameter and $\{a, b, c\} \in \mathcal{R}^+$ are hyper parameters. Given some observations **x**, we multiply its likelihood under the Gamma distribution by the prior on shape (7) to obtain an expression for the posterior distribution $q(\alpha)$

$$q(\alpha) \propto \frac{\hat{a}^{\alpha-1} R^{\alpha \hat{c}}}{\Gamma(\alpha)^{\hat{b}}} \tag{8}$$

with

$$\hat{a} = a \prod_{i=1}^{n} x_i, \qquad \hat{b} = b + n, \qquad \hat{c} = c + n. \tag{9}$$

Finding the posterior expectation of α, $\hat{\alpha}$, implies computing the expectation of $p(\alpha)$. Here we use the Laplace approximation to (7), which can be shown to be a Gaussian with mean

$$m = \Psi^{-1}\Big(\frac{\log a + c \log R}{b}\Big),$$

and precision $b\Psi_1(m)$, where $\Psi_1(\alpha) = \frac{d\Psi(\alpha)}{d\alpha}$. Consequently we approximate the posterior expected shape by

$$\hat{\alpha} \approx \Psi^{-1}\Big(\frac{\log \hat{a} + \hat{c} \log R}{\hat{b}}\Big). \tag{10}$$

We note here that the expectation of the Laplace approximation to $q(\alpha)$ corresponds to the maximum a posteriori (MAP) estimate of $q(\alpha)$. The use of the Laplace approximation in this context then reduces to using a MAP estimation in place of the expected value. Note also that in order to compute the expectation (10) we need to estimate $\log \hat{a}$

$$\log \hat{a} = \log a + \sum_{i=1}^{n} \log x_i, \tag{11}$$

Algorithm 1 BL1

Require: $\mathbf{x} = \{x_1, \ldots, x_n\}, x_i > 0$ and $\{a, b, c, d, e\}$

$\mu = \frac{1}{n}\sum_{i=1}^{n} x_i$

$v = \frac{1}{n-1}\sum_{i=1}^{n}(x_i - \mu)^2$

$\alpha = \frac{\mu^2}{v}$

$\hat{e} = e + \sum_{i=1}^{n} x_i$

$\log \hat{a} = \log a + \sum_{i=1}^{n} \log x_i$

$\hat{b} = b + n$

$\hat{c} = c + n$

repeat

$\alpha \leftarrow \Psi^{-1}\Big(\frac{\log \hat{a} + \hat{c}\big(\log(d+n\alpha) - \log(\hat{e})\big)}{\hat{b}}\Big)$

until convergence

$\hat{\alpha} = \alpha$

$\hat{d} = d + n\hat{\alpha}$

$\hat{\beta} = \frac{\hat{d}}{\hat{e}}$

return $\hat{\alpha}$,$\hat{\beta}$

and not longer require $\hat{a}$ in equation (9). This fact has the advantage of avoiding numerical issues for large sample sizes. Further, substituting $\hat{R} = \frac{\hat{d}}{\hat{e}} = \frac{d+n\alpha}{e+\sum_{i=1}^{n} x_i}$ into equation (16) we obtain an expression with no R dependence which is more compact for algorithmic use, namely

$$\hat{\alpha} \approx \Psi^{-1}\Big(\frac{\log \hat{a} + \hat{c}\big(\log(d+n\alpha) - \log(e+\sum_{i=1}^{n} x_i)\big)}{\hat{b}}\Big). \tag{12}$$

Algorithm 1 summarizes this first Bayesian approach for learning the Gamma parameters which we denote as $BL1$. First, the MM algorithm is used to initialize α (equation (23) left panel); $\hat{e}$, $\log \hat{a}$, $\hat{b}$ and $\hat{c}$ are computed through equations (5 right side), (11) and (9). Then equation (12) is iterated till convergence to obtain the expected $\hat{\alpha}$. Finally $\hat{d}$ is computed using equation (5 left side) and $\hat{\beta}$ through (6).

2.3 Likelihood approximation and its conjugate prior on α (BL2)

Given an initial parameter value estimation α, we use an approximation to the log-likelihood as presented in [8], namely

$$\log \mathcal{G}(\mathbf{x}|\alpha) = k_0 + k_1\alpha + k_2 \log \alpha. \tag{13}$$

Learning the parameters k_0, k_1, k_2 is straightforward (check Appendix B). We then define a prior on α that is conjugate with this approximated log-likelihood, namely

$$\log p(\alpha) \propto w_0 + w_1\alpha + w_2 \log \alpha \tag{14}$$

for any set of hyper parameter values w_0, w_1, w_2. It is straightforward to observe that the posterior values for the hyper parameters are

$$\tilde{w}_i = w_i + k_i, \qquad \forall i \in \{0, 1, 2\}. \tag{15}$$

Once more we use the Laplace approximation to (14), which can be shown to be a Gaussian with mean

$$m = \frac{-\tilde{w_2}}{\tilde{w_1}}$$

and precision $\frac{\tilde{w_2}}{\alpha^2}$. Consequently we approximate the posterior expected shape by

$$\hat{\alpha} \approx \frac{-\tilde{w_2}}{\tilde{w_1}}. \tag{16}$$

Again, the Laplace approximation corresponds to the maximum a posteriori (MAP) estimate. Algorithm 2 summarizes this alternative for Bayesian learning the parameters of the Gamma distribution which will be further denoted as $BL2$.

Algorithm 2 BL2

Require: $\mathbf{x} = \{x_1, \ldots, x_n\}, x_i > 0$ and $\{w_1, w_2\}$
  $\mu = \frac{1}{n}\sum_{i=1}^{n} x_i$
  $v = \frac{1}{n-1}\sum_{i=1}^{n}(x_i - \mu)^2$
  $\alpha = \frac{\mu^2}{v}$
  repeat
    $k_1 = n\big(\overline{\log \mathbf{x}} - \Psi(\alpha) - \log \overline{\mathbf{x}} + \log \alpha - \alpha\Psi'(\alpha) + 1\big)$
    $k_2 = n\alpha^2\Psi'(\alpha) - n\alpha$
    $\tilde{w_1} = w_1 + k_1$
    $\tilde{w_2} = w_2 + k_2$
    $\alpha \leftarrow \frac{-\tilde{w_2}}{\tilde{w1}}$
  until convergence
  $\hat{\alpha} = \alpha$
  $\hat{d} = d + n\hat{\alpha}$
  $\hat{\beta} = \frac{\hat{d}}{\hat{e}}$
  return $\hat{\alpha}$,$\hat{\beta}$

3 Results

In section 3.1 we highlight the relationship between the ML algorithms presented in [8] and the Bayesian ones presented here. Both ML algorithms (ML1 and ML2) can be found in Appendix B. In section 3.2 we present numerical results obtained using synthetic data and in section 3.3 we consider the hyper parameter value setting for the Bayesian shape priors.

3.1 Algorithms comparison

It is to note that ML and BL algorithms have notable similarities; first, note that the ML and the Bayesian scale estimators presented are

$$\hat{\beta}_{ML} = \frac{\mu}{\alpha} = \frac{\sum_{i=1}^{n} x_i}{n\alpha} \qquad \hat{\beta}_{BL} = \hat{R}^{-1} = \frac{\hat{e}}{\hat{d}} = \frac{e + \sum_{i=1}^{n} x_i}{d + n\alpha} \tag{17}$$

respectively. It is easy to observe that

$$\lim_{d=e\to 0} \hat{\beta}_{BL} \to \hat{\beta}_{ML}, \tag{18}$$

which means that the Bayesian scale estimator tends to the ML one in the limit of an infinite variance gamma prior over β. With respect to the shape parameter α, since initializing $b = c$ results in $\hat{b} = \hat{c}$ (in BL1), and taking again the same limit, we can rewrite the $BL1$ α update as

$$\lim_{\substack{d=e\to 0 \\ b=c}} \hat{\alpha}_{BL} \to \Psi^{-1}\Big(\frac{\log \hat{a}}{\hat{b}} + \log(n\alpha) - \log \sum_{i=1}^{n} x_i\Big) \tag{19}$$

Comparing this to the ML α update in the case of a linear constrain on α (ML1)

$$\hat{\alpha}_{ML1} = \Psi^{-1}(\log \alpha + \overline{\log \mathbf{x}} - \log \bar{x}) = \Psi^{-1}(\overline{\log \mathbf{x}} + \log n\alpha - \log \sum_{i=1}^{n} x_i) \tag{20}$$

it is clear that both models perform an iterative update of the shape parameter dependent in the inverse of the gamma function of $\log n\alpha$ plus a data dependent constant which differs for ML1 and BL1; they also have the same dependence on one of the data dependent terms, $-\log \sum_{i=1}^{n} x_i$, independently of the hyper parameter values used for the prior on α (under the assumption b=c). Further, since $\log \hat{a} = \log a + \sum_{i=1}^{n} \log x_i$ and $\hat{b} = b + n$, in the extreme case of considering a very small hyper parameter value for b, we can simplify the BL1 estimation update further, resulting in

$$\lim_{\substack{d=e\to 0 \\ b=c\to 0}} \hat{\alpha}_{BL} \to \Psi^{-1}\Big(\frac{\log a}{n} + \overline{\log \mathbf{x}} + \log(n\alpha) - \log \sum_{i=1}^{n} x_i\Big) \tag{21}$$

Remembering that the posterior estimation for b and c are $\hat{b} = b + n$ and $\hat{c} = c + n$, this is a very interesting result that shows that when choosing little informative hyper parameter values for b and c, the BL1 estimation involves a small sample bias correction term $\frac{-\log a}{n}$ which tends to 0 as the number of observed samples increases. Further, there is, by construction, a close relationship between ML2 and BL2; in fact they are equivalent when considering a flat prior over α, that is when $w_1 = w_2 = 0$ in equation (14).

3.2 Numerical results

In this section we perform a numerical comparison between the two maximum likelihood (ML) algorithms presented in [9] and the Bayesian approaches (BL) presented in this work. For each simulation we generate varying amount of N samples from a Gamma distribution with fixed parameters α and β and we applied the four considered algorithms. The parameter values α and β are initialized to different random positive numbers at each simulation. For each N we performed 500 different simulations. The ML and the BL algorithms are considered to converge when the relative α parameter change between two consecutive iterations is smaller than $10^{-}6$. In this example we use the information from the previous section to fix the hyper parameters to values for which the BL solutions tend to the ML ones; we used the hyper parameter values $d = e = 0.001$, $a = 1$ and $b = c = 0.001$ for BL1 and $w_0 = 1$, $w_1 = w_2 = 0$ for BL2. To asses the quality of the estimators

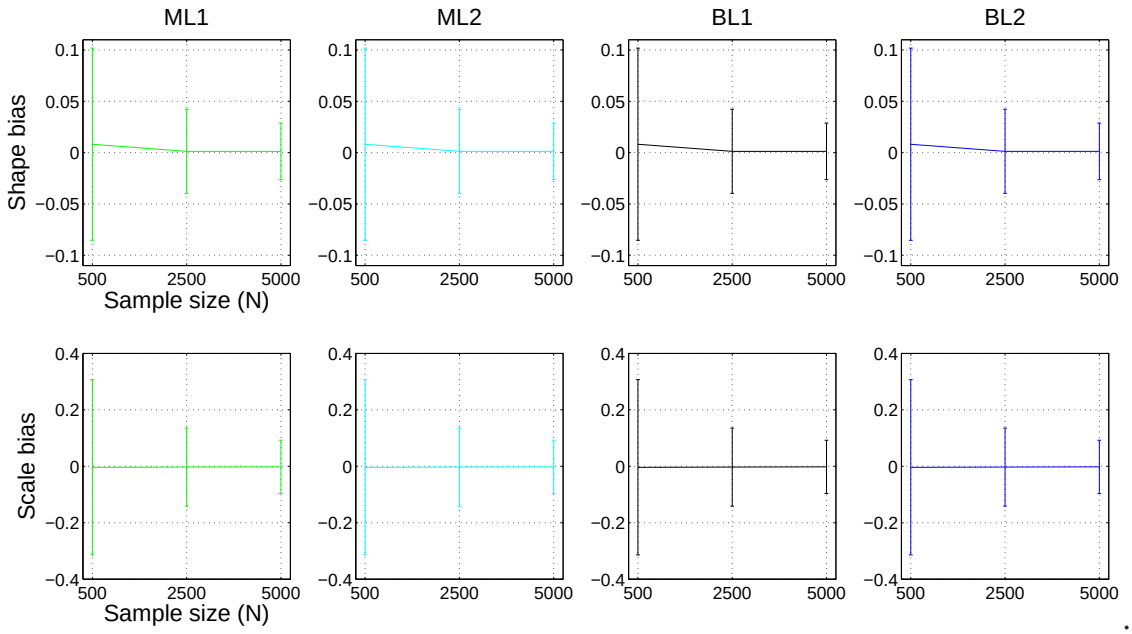

Figure 1: .

we computed KL-divergences between the true distribution and the estimated ones. Then, for each sample size (N), we performed a paired t-test between the KL-divergences of each possible pair of algorithms and found no statistically significant differences between them. This is not a surprising result since both ML algorithms converge to the same solution and the priors in the Bayesian setting are chosen to tend to obtain a estimator tending to the ML one. In each sub-figure of figure 1 we present error bars (mean and standard deviation) of the bias in the estimation of each parameter, shape in first row, scale in the second one. Each algorithm is presented in a different column and the x-axis represent the number of samples used for the estimation. These results are, for the ML case, in agreement with the ones presented in [5]; the ML shape estimators are upwards biased for the shape and downwards biased for the scale parameter. The bias of the presented Bayesian algorithms is equivalent to that of the existing ML algorithms.

Figure 2 shows boxplots of the computational costs in seconds; the mean, 25^{-th} and 75^{-th} percentile are presented on each box and the red crosses represent outliers of the estimated distributions. The cost of ML1 and BL1 is equivalent while ML2 and BL2 are much faster. It is to note that in all cases the iterative process depends only on the inverse of the digamma function of a real number and there is no algorithmic cost N dependency in the loop.

3.3 Shape conjugate prior hyper parameters values

It is easy to see flat prior for the shape parameter in the BL2 case can be obtained by choosing $w_1 = w_2 = 0$ in equation (14). On the other side, looking at equation (7) it is clear that choosing a flat prior for BL1 is not straightforward. As we showed in section 3.3, choosing $b = c \to 0^+$ and $a \in \mathcal{R}^+$ results in a prior providing a small bias correction which tends to zero for high sample sizes. To get some more intuition we next provide some numerical examples illustrating such prior for different hyper parameter values and demonstrating the behavior of the posterior hyper

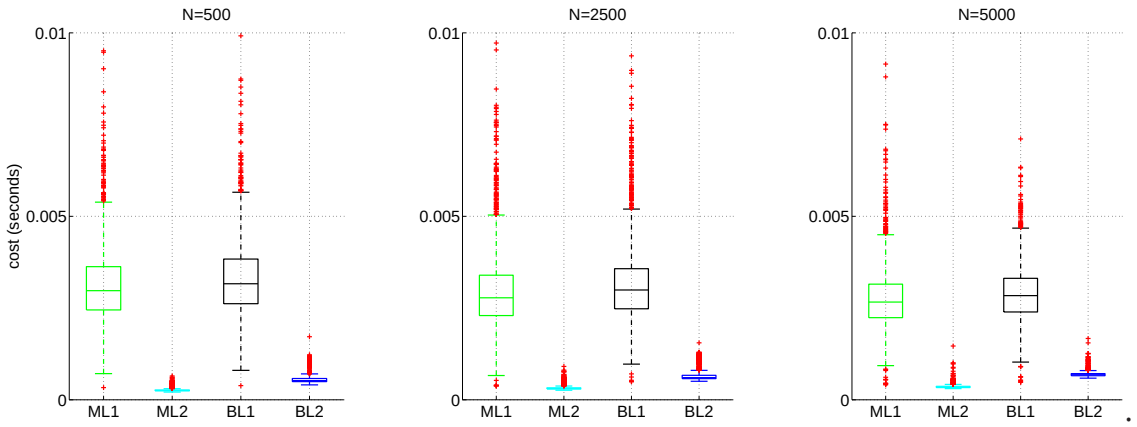

Figure 2: Boxplots of the computational time required by the different algorithms expressed in seconds. Each column represent a different sample size.

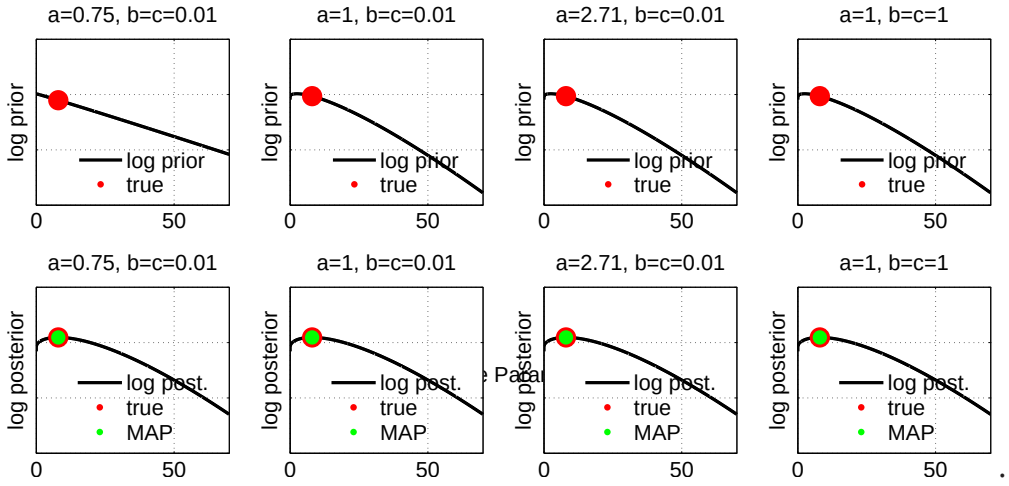

Figure 3: Top row shows the log-prior $\log p(\alpha|a,b,c,\beta)$ and the bottom row the log-posterior $\log p(\alpha|\hat{a},\hat{b},\hat{c},\hat{\beta})$. The red dots mark the shape parameter α from which data was generated. The green dots represent the Bayesian posterior estimation.

parameter updates (equations 9,11) and of the Laplace approximation (equation 10) used to obtain the expected shape value $\hat{\alpha}$.

We generate 1000 samples from an Gamma distribution with parameter values $\alpha = 10$ and $\beta = 25$. $\hat{\beta}$ is computed through equation (23) left panel followed by equations (5,6) where a flat prior was used, namely $d = e = 0.01$. For the shape prior hyper parameters we consider different values a, b, c, and $\hat{\alpha}$ is computed independently on each case through equations (9,11,16). In Figure 3, the top row shows the log-prior $\log p(\alpha|a,b,c,\hat{\beta})$ for the parameter values a, b, c indicated in the titles as a function of the shape value α represented in the x-axis. The bottom row shows the corresponding log-posteriors $\log p(\alpha|\hat{a},\hat{b},\hat{c},\hat{\beta})$. In all cases the red dot represents the true shape value, $\alpha = 10$, and the green ones in the bottom row represents the posterior estimated value $\hat{\alpha}$. The expectation we obtained for the posterior of α, $\hat{\alpha}$ marked as as green circle, is a very good approximation. Since the parameter value for the scale β was also estimated using Bayesian

inference and its part of the prior and posterior on α, the figure indirectly confirms also the proper behaviour of the scale Bayesian estimation procedure.

These examples are representative of the general behaviour observed for a broad range of hyper parameter values $\{a, b, c\}$. Although this shows the weak dependence on the prior when estimating the posterior, further analyses relating the effect of the hyper parameter value of a with relation to the well known bias correction strategies for the parameters of the Gamma distribution [5] would be of interest and are focus of ongoing research.

4 Discussion

We reviewed two maximum likelihood algorithms for the estimation of the parameters of a Gamma distribution and developed two Bayesian algorithms; the first one uses an unnormalized conjugate prior to the shape parameter and the second one uses a likelihood approximation and prior conjugated with the approximation. In both cases we use the Laplace approximation to compute the required expectations. We performed a theoretical analyses which showed which hyper-parameter values for the conjugate priors will recover the ML algorithms and we also show numerically that with such hyper-parameter choices the new Bayesian algorithms converge towards the ML solutions. We also evaluated numerically the bias in the parameters estimation and show that the Bayesian algorithms have the same bias properties as the ML ones. Including bias corrections in the Bayesian setting for small sample sizes could be achieved as showed in [5] for the ML case and is part of ongoing research. The Bayesian Learning approaches introduced in section 2 has important implications since it can be used inside more complex Bayesian inference tasks, such as fitting mixture models involving Gamma distributions within a (Variational) Bayesian framework to be used for example in the context of medical image segmentation [10]. Extending this beyond modeling one dimensional data, the presented methodology can also be included in multivariate models, e.g. for Bayesian ICA or Non-negative matrix factorizations.

References

[1] H. Aksoy. Use of gamma distribution in hydrological analysis. *Journal of Engineering Enviromental Science*, (24):419–228, 2000.

[2] C. M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer, 1 edition, 2007.

[3] J. D. Cook. Inverse gamma distribution. *Online: http://www. johndcook. com/inverse gamma. pdf, Technical. Report*, 2008.

[4] D. Fink. A compendium of conjugate priors, 1997.

[5] D. Giles and H. Feng. Bias of the maximum likelihood estimators of the two-parameter gamma distribution revisited. Econometrics Working Papers 0908, Department of Economics, University of Victoria, 2009.

[6] A. Khalili, D. Potter, P. Yan, L. Li, J. Gray, T. Huang, and S. Lin. Gamma-normal-gamma mixture model for detecting differentially methylated loci in three breast cancer cell lines. *Cancer Inform*, 3:43–54, 2007.

[7] D. J. C. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambrdige University Press, 2003.

[8] T. P. Minka. Beyond Newton's method. Aug. 2000.

[9] T. P. Minka. Estimating a gamma distribution. *Microsoft Research*, 2002.

[10] M. Woolrich, T. Behrens, C. Beckmann, and S. Smith. Mixture models with adaptive spatial regularization for segmentation with an application to fmri data. *Medical Imaging, IEEE Transactions on*, 24(1):1–11, Jan 2005.

[11] S. Yue. A bivariate gamma distribution for use in multivariate flood frequency analysis. *Hydrological Processes*, 15(6):1033–1045, 2001.

5 Appendices

A Method of Moments (MM)

The first two moments of a Gamma distribution are [3]

$$\mathbb{E}_{\mathcal{G}}[x] = \alpha\beta, \qquad \mathbb{E}_{\mathcal{G}}[(x - \mathbb{E}_{\mathcal{G}}[x])^2] = \alpha\beta^2. \tag{22}$$

Using a Gaussian approximation to equation (22)

$$\mu \approx \alpha\beta, \qquad v \approx \alpha\beta^2,$$

where μ and v are the mean and variance estimated from the observed data vector $\mathbf{x} = \{x_1, \ldots, x_n\}$, and solving the linear system for α and β we obtain the MM estimation for the Gamma distribution parameters, namely

$$\hat{\alpha} \approx \frac{\mu^2}{v}, \qquad \hat{\beta} \approx \frac{v}{\mu} \tag{23}$$

Algorithm 3 summarizes the process.

Algorithm 3 MM for Gamma

Require: $\mathbf{x} = \{x_1, \ldots, x_n\}, x_i > 0$
$\mu = \frac{1}{n}\sum_{i=1}^{n} x_i$
$v = \frac{1}{n-1}\sum_{i=1}^{n}(x_i - \mu)^2$
$\hat{\alpha} = \frac{\mu^2}{v}$
$\hat{\beta} = \frac{v}{\mu}$
return $\hat{\alpha}$,$\hat{\beta}$

Obviously this approximation is very fast since it simply requires estimation of mean and variance. Nevertheless it will provide more biased estimations as the underlying Gamma distribution deviates more from a Gaussian i.e. as the Gamma distribution higher order moments more deviate from zero.

B Maximum Likelihood (ML)

In this subsection we describe the two algorithms for ML learning of Gamma distributions parameters originally presented in [9]. The log-likelihood of the positive vector of observations $\mathbf{x} = \{x_1, \ldots, x_n\}$ under the Gamma distribution (1) can be written as

$$\log \mathcal{G}(\mathbf{x}|\alpha, \beta) = n(\alpha - 1)\overline{\log \mathbf{x}} - n \log \Gamma(\alpha) - n\alpha \log \beta - \frac{n\overline{\mathbf{x}}}{\beta} \tag{24}$$

where for easy of notation the upper bar operand denotes the arithmetic mean operator. Finding ML estimations for the parameters is achieved by maximizing equation (24) with respect to the parameters $\{\alpha, \beta\}$; it is easy to verify that equation (24) has a maximum at

$$\beta = \frac{\overline{\mathbf{x}}}{\alpha}, \tag{25}$$

and that direct maximization of equation (24) with respect to α it is not possible. Substituting equation (25) into equation (24) gives

$$\log \mathcal{G}(\mathbf{x}|\alpha) = n(\alpha - 1)\overline{\log \mathbf{x}} - n \log \Gamma(\alpha) - n\alpha \log \overline{\mathbf{x}} + n\alpha \log \alpha - n\alpha \tag{26}$$

Direct maximization of equation (26) with respect to α is (obviously) also not possible. Using a linear constrain around a given value for α, denoted as α_0, we have

$$\alpha \log(\alpha) \geq (1 + \log \alpha_0)(\alpha - \alpha_0) + \alpha_0 \log(\alpha_0). \tag{27}$$

Substituting (27) into equation (26) provides a lower bound on the log-likelihood. Differentiating with respect to α, equaling to zero and solving for α we have that

$$\hat{\alpha} = \Psi^{-1}(\log \alpha_0 + \overline{\log \mathbf{x}} - \log \bar{x}) \tag{28}$$

where Ψ represents the *digamma function*. For an initial value of α equation (28) can be iterated till convergence to obtain the desired value $\hat{\alpha}$ and then $\hat{\beta}$ can be computed using equation (25). In this work we initialize the value of α using the method of moments and Algorithm 4 summarizes the process.

Algorithm 4 ML1 (linear constrain)

Require: $\mathbf{x} = \{x_1, \ldots, x_n\}, x_i > 0$
  $\mu = \bar{x} = \frac{1}{n}\sum_{i=1}^{n} x_i$
  $v = \frac{1}{n-1}\sum_{i=1}^{n}(x_i - \mu)^2$
  $\alpha = \frac{\mu^2}{v}$
  repeat
    $\alpha \leftarrow \Psi^{-1}(\log \alpha + \overline{\log \mathbf{x}} - \log \bar{x})$
  until convergence
  $\hat{\alpha} = \alpha$
  $\hat{\beta} = \frac{\bar{x}}{\hat{\alpha}}$
  return $\hat{\alpha}$,$\hat{\beta}$

Algorithm 4 provides an accurate ML solution but it suffers from slow convergence. To accelerate the process, one can approximate Gamma log-likelihood presented in equation (26) by

$$f(\alpha) = k_0 + k_1\alpha + k_2 \log \alpha. \tag{29}$$

Taking first and second derivatives of $f(\alpha)$

$$f'(\alpha) = k_1 + \frac{k_2}{\alpha}, \qquad f''(\alpha) = -\frac{k_2}{\alpha^2}, \tag{30}$$

and matching $f(\alpha)$ and its derivatives to $\log \mathcal{G}(\mathbf{x}|\alpha)$ and its first two derivatives with respect to α respectively, we obtain values for k_0, k_1, k_2. If $f''(\alpha) < 0$, $f(\alpha)$ has a maximum at $\hat{\alpha}$ iff $f'(\hat{\alpha}) = 0$. This process gives us the update rule

$$\frac{1}{\alpha} \leftarrow \frac{1}{\alpha} + \frac{\overline{\log x} - \log \bar{x} + \log \alpha - \Psi(\alpha)}{\alpha^2(\frac{1}{\alpha} - \Psi'(\alpha))} \tag{31}$$

Algorithm 5 ML2 (non-linear approximation)

Require: $\mathbf{x} = \{x_1, \ldots, x_n\}, x_i > 0$
$\mu = \bar{x} = \frac{1}{n}\sum_{i=1}^{n} x_i$
$v = \frac{1}{n-1}\sum_{i=1}^{n}(x_i - \mu)^2$
$\alpha = \frac{\mu^2}{v}$
repeat
$\frac{1}{\alpha} \leftarrow \frac{1}{\alpha} + \frac{\overline{\log x} - \log \bar{x} + \log \alpha - \Psi(\alpha)}{\alpha^2(\frac{1}{\alpha} - \Psi'(\alpha))}$
until convergence
$\hat{\alpha} = \alpha$
$\hat{\beta} = \frac{\bar{x}}{\hat{\alpha}}$
return $\hat{\alpha}$,$\hat{\beta}$

As before we use the MM relationship for the shape parameter (equation 23, left side) to initialize α. After iteration till convergence to get $\hat{\alpha}$ we use equation (25) to get $\hat{\beta}$. Algorithm 5 summarizes the process:

The two ML algorithms presented here deviate from the originally presented in [9] in the initialization. While we use the MM for initialization, in the original work a closed form estimation was used.